

IN THE CLAIMS:

1 1. (ORIGINAL) A method for enabling recovery from two or fewer concurrent failures
2 of storage devices in a storage array, the method comprising the steps of:

3 providing the array with a predetermined number of storage devices, including a
4 plurality of first devices configured to store data and row parity, and one diagonal parity
5 device configured to store diagonal parity, wherein the predetermined number of storage
6 devices n is $p+1$ and wherein p is a prime number;

7 dividing each device into blocks;

8 organizing the blocks into stripes that contain a same number of blocks in each
9 device, wherein each stripe comprises $n-2$ rows of blocks;

10 defining the diagonal parity along diagonal parity sets that span the first devices,
11 wherein the diagonal parity sets wrap around within a group of $n-2$ rows so that all blocks
12 belonging to diagonal parity sets of a stripe are stored in the stripe; and

13 computing and storing the diagonal parity for all of the diagonal parity sets except
14 one on the diagonal parity device.

1 2. (ORIGINAL) The method of Claim 1 wherein row parity blocks in a stripe are all
2 stored on a single device.

1 3. (ORIGINAL) The method of Claim 1 wherein the first devices include a plurality of
2 data devices, and wherein the data devices are not all present and absent data devices are
3 treated as containing zero-valued data.

1 4. (ORIGINAL) The method of Claim 1 wherein computation of row parity blocks is
2 replaced by a redundant encoding of data blocks in a stripe that is tolerant of single de-
3 vice failures.

1 5. (ORIGINAL) The method of Claim 1 wherein a total length of the devices is unequal
2 and wherein each row has one row parity block and one diagonal parity block.

1 6. (ORIGINAL) The method of Claim 1 wherein locations of parity blocks shift from
2 device to device within different stripes.

1 7. (ORIGINAL) A system configured to enable recovery from two or fewer concurrent
2 failures of two storage devices, the system comprising:

3 an array having a predetermined number of storage devices, including a plurality
4 of first devices configured to store data and row parity, and one diagonal parity device
5 configured to store diagonal parity, wherein the predetermined number of storage devices
6 n is $p+1$ and wherein p is a prime number;

7 a storage operating system including a device storage layer configured to imple-
8 ment a row-diagonal (R-D) parity technique that computes the diagonal parity along di-
9 agonal parity sets that span the first devices, and that stores the diagonal parity for all of
10 the diagonal parity sets except one on the diagonal parity device; and

11 a processing element configured to execute the storage operating system to
12 thereby invoke storage access operations to and from the array in accordance with the R-
13 D parity technique.

- 1 8. (ORIGINAL) The system of Claim 7 wherein row parity blocks in a stripe are all
2 stored on a single device.
- 1 9. (ORIGINAL) The system of Claim 7 wherein computation of row parity blocks is re-
2 placed by a redundant encoding of data blocks in a stripe that is tolerant of single device
3 failures.
- 1 10. (ORIGINAL) The system of Claim 7 wherein the device storage layer is a RAID
2 layer and wherein the storage devices are disks.
- 1 11. (ORIGINAL) The system of Claim 10 wherein the RAID layer further divides each
2 disk into blocks and organizes the blocks into stripes.
- 1 12. (ORIGINAL) The system of Claim 11 wherein each stripe comprises $n-2$ rows of
2 blocks, with $n=p+1$, wherein each row contains one block from each disk.
- 1 13. (ORIGINAL) The system of Claim 12 wherein a total length of the devices is un-
2 equal and wherein each row has one row parity block and one diagonal parity block.
- 1 14. (ORIGINAL) The system of Claim 10 wherein the RAID layer is implemented in
2 logic circuitry.

- 1 15. (ORIGINAL) The system of Claim 7 wherein the devices are one of video tape, magnetic tape, optical, DVD, bubble memory, magnetic disk, electronic random access memory and micro-electro mechanical storage devices.
- 1 16. (ORIGINAL) The system of Claim 7 wherein the devices are media adapted to store information contained within the data and parity blocks.
- 1 17. (ORIGINAL) The system of Claim 7 wherein the first devices include a plurality of data devices, and wherein the data devices are not all present and absent data devices are treated as containing zero-valued data.
- 1 18. (ORIGINAL) Apparatus for enabling recovery from two or fewer concurrent failures of two storage devices in a storage array, the apparatus comprising:
 - 3 means for providing the array with a predetermined number of storage devices, including a plurality of first devices configured to store data and row parity, and one diagonal parity device configured to store diagonal parity, wherein the predetermined number of storage devices n is $p+1$ and wherein p is a prime number;
 - 7 means for computing the diagonal parity along diagonal parity sets that span the first devices; and
 - 9 means for storing the diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

- 1 19. (ORIGINAL) The apparatus of Claim 18 wherein row parity blocks in a stripe are all
- 2 stored on a single device.

- 1 20. (ORIGINAL) The system of Claim 18 wherein computation of row parity blocks is
- 2 replaced by a redundant encoding of data blocks in a stripe that is tolerant of single de-
- 3 vice failures.

- 1 21. (ORIGINAL) The apparatus of Claim 18 further comprising:
 - 2 means for dividing each device into blocks; and
 - 3 means for organizing the blocks into stripes.

- 1 22. (ORIGINAL) The apparatus of Claim 21 wherein the means for storing comprises
- 2 means for storing diagonal parity blocks on the diagonal parity disk for each of the di-
- 3 agonal parity sets of a stripe except one.

- 1 23. (ORIGINAL) The apparatus of Claim 18 wherein the first devices include a plurality
- 2 of data devices, and wherein the data devices are not all present and absent data devices
- 3 are treated as containing zero-valued data.

- 1 24. (ORIGINAL) A computer readable medium containing executable program instruc-
- 2 tions for enabling recovery from two or fewer concurrent failures of two storage devices

3 in a storage array, the executable program instructions comprising program instructions
4 for:

5 providing the array with a predetermined number of storage devices, including a
6 plurality of first devices configured to store data and row parity, and one diagonal parity
7 device configured to store diagonal parity, wherein the predetermined number of storage
8 devices n is $p+1$ and wherein p is a prime number;

9 computing the diagonal parity along diagonal parity sets that span the first de-
10 vices; and

11 storing the diagonal parity for all of the diagonal parity sets except one on the di-
12 agonal parity device.

1 25. (ORIGINAL) The computer readable medium of Claim 24 wherein row parity blocks
2 in a stripe are all stored on a single device.

1 26. (ORIGINAL) The computer readable medium of Claim 24 wherein computation of
2 row parity blocks is replaced by a redundant encoding of data blocks in a stripe that is
3 tolerant of single device failures.

1 27. (ORIGINAL) The computer readable medium of Claim 24 further comprising pro-
2 gram instructions for:

3 dividing each device into blocks; and

4 organizing the blocks into stripes.

1 28. (ORIGINAL) The computer readable medium of Claim 27 wherein the program in-
2 struction for storing comprises program instructions for storing diagonal parity blocks on
3 the diagonal parity disk for each of the diagonal parity sets of a stripe except one.

1 29. (ORIGINAL) The computer readable medium of Claim 24 wherein the first devices
2 include a plurality of data devices, and wherein the data devices are not all present and
3 absent data devices are treated as containing zero-valued data.

1 30. (ORIGINAL) A method for enabling recovery from two or fewer concurrent failures
2 of two storage devices in a storage array, the method comprising the steps of:

3 providing the array with a predetermined number of storage devices equal to n ,
4 where $n=p+1$ and p is a prime number, including $n-2$ data devices configured to hold
5 data, one row parity device and one diagonal parity device;

6 dividing each device into blocks of fixed size;

7 organizing the blocks into stripes, wherein each stripe comprises $n-2$ rows of
8 blocks such that each device contains $n-2$ blocks per stripe;

9 computing the diagonal parity along $n-2$ of $n-1$ parity sets (diagonal) of blocks
10 that span the data devices and the row parity device, wherein the blocks of a diagonal are
11 contained within a stripe, each diagonal containing a set of $n-2$ blocks selected from
12 among data and row parity blocks in a stripe such that no diagonal contains more than
13 one block from a same row, each diagonal further containing $n-2$ blocks from a group of
14 $n-2$ rows and no diagonal containing two blocks from a same device;

15 storing, on the row parity device, row parity for rows of blocks computed across
16 all of the data devices such that each row contains one block from each data device and
17 no row contains two blocks from a same diagonal; and

18 storing, in the $n-2$ blocks in the stripe on the diagonal parity device, the diagonal
19 parity of $n-2$ of the $n-1$ diagonals, wherein selection of the diagonals stored on the diag-
20 onal parity device is predetermined and wherein an order of placement of the $n-2$ diagonal
21 parity blocks in the stripe on the diagonal parity device is predetermined.

1 31. (ORIGINAL) The method of Claim 30 wherein the rows containing each stripe are
2 contiguous.

1 32. (ORIGINAL) The method of Claim 31 wherein the diagonals are formed by number-
2 ing the data, row parity and diagonal parity devices from 0 to $n-1$, and numbering the
3 rows from 0 to $n-3$, then assigning the block at device i , row j to diagonal $(i+j) \bmod(n-2)$.

1

1 33. (ORIGINAL) The method of Claim 30 wherein the data devices are not all present
2 and wherein absent data devices are treated as containing zero-valued data.

1 34. (ORIGINAL) The method of Claim 30 wherein a total length of the devices is un-
2 equal and wherein each row has one row parity block and one diagonal parity block.

1 35. (ORIGINAL) The method of Claim 30 wherein the storage devices are one of video
2 tape, magnetic tape, optical, DVD, bubble memory, electronic random access memory,

3 micro-electro mechanical and magnetic disk devices, and wherein the magnetic disk de-
4 vices are one of hard disk drives and direct access storage devices.

1 36. (ORIGINAL) A method for enabling recovery from two or fewer concurrent failures
2 of two storage devices in a storage array, the method comprising the steps of:

3 providing the array with a predetermined number of storage devices, including a
4 plurality of first devices configured to store data and row parity, and one diagonal parity
5 device configured to store diagonal parity;

6 computing the diagonal parity along diagonal parity sets that span the first de-
7 vices; and

8 storing the diagonal parity for all of the diagonal parity sets except one on the di-
9 agonal parity device.

1 37. (ORIGINAL) The method of Claim 36 further comprising the steps of:

2 dividing each device into blocks; and

3 organizing the blocks into stripes.

1 38. (ORIGINAL) The method of Claim 37 wherein the step of storing comprises the step
2 of storing diagonal parity blocks on the diagonal parity disk for each of the diagonal par-
3 ity sets of a stripe except one.

1 39. (ORIGINAL) The method of Claim 36 wherein the first devices include a plurality of
2 data devices, and wherein the data devices are not all present and absent data devices are
3 treated as containing zero-valued data.

1 40. (ORIGINAL) The method of Claim 36 wherein the predetermined number of storage
2 devices is n , and wherein $n=p+1$ and p is a prime number.

1 41. (ORIGINAL) A method for enabling recovery from two or fewer concurrent failures
2 of two storage devices in a storage array, the method comprising the steps of:

3 providing the array with a predetermined number of storage devices equal to n ,
4 where $n=p+1$ and p is a prime number, including $n-2$ data devices configured to hold
5 data, one row parity device and one diagonal parity device;

6 dividing each device into blocks of fixed size;

7 organizing the blocks into stripes, wherein each stripe comprises $n-2$ rows of
8 blocks such that each device contains $n-2$ blocks per stripe and wherein locations of par-
9 ity blocks shift from device to device within different stripes;

10 computing the diagonal parity along $n-2$ of $n-1$ parity sets (diagonal) of blocks
11 that span the data devices and the row parity device, wherein the blocks of a diagonal are
12 contained within a stripe, each diagonal containing a set of $n-2$ blocks selected from
13 among data and row parity blocks in a stripe such that no diagonal contains more than
14 one block from a same row, each diagonal further containing $n-2$ blocks from a group of
15 $n-2$ rows and no diagonal containing two blocks from a same device;

16 storing, on the row parity device, row parity for rows of blocks computed across
17 all of the data devices such that each row contains one block from each data device and
18 no row contains two blocks from a same diagonal; and

19 storing, in the $n-2$ blocks in the stripe on the diagonal parity device, the diagonal
20 parity of $n-2$ of the $n-1$ diagonals, wherein selection of the diagonals stored on the diag-
21 onal parity device is predetermined and wherein an order of placement of the $n-2$ diagonal
22 parity blocks in the stripe on the diagonal parity device is predetermined.

1 42. (PREVIOUSLY PRESENTED) A method for enabling recovery from two or fewer
2 concurrent failures of storage devices in a storage array, comprising:

3 storing data and row parity in stripes on a plurality of first storage devices, the
4 stripes configured as rows of blocks on the first devices;

5 defining diagonal parity sets that span the first devices;

6 computing a diagonal parity for the diagonal parity sets; and

7 storing the diagonal parity on a diagonal parity device, the diagonal parity device
8 separate from the first storage devices, the storage array including the plurality of first
9 storage devices and the diagonal parity device.

1 43. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:

2 predetermining the a number of storage devices as n , where n is $p+1$, and where p
3 is a prime number.

- 1 44. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 storing row parity blocks on a single storage device of the plurality of first storage
 - 3 devices.
- 1 45. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 configuring the plurality of first storage devices so that not all first storage de-
 - 3 vices are present, and treating absent data storage devices as containing zero valued data.
- 1 46. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 computing row parity blocks by an encoding of data blocks that is tolerant of sin-
 - 3 gle device failures.
- 1 47. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 configuring each row to have one row parity block and one diagonal parity block.
- 1 48. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 placing parity blocks in different locations in different stripes from device to de-
 - 3 vice.

1 49. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 implementing a row-diagonal (R-D) parity technique that computes the diagonal
3 parity along diagonal parity sets that span the first devices, and that stores the diagonal
4 parity for all of the diagonal parity sets except one on the diagonal parity device.

1 50. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 invoking storage access operations to and from the array in accordance with a R-
3 D parity technique.

1 51. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 storing row parity blocks computed for a stripe on a single first storage device.

1 52. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 configuring the array as a RAID system of computer storage units.

1 53. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 configuring the array as a RAID system of computer storage disks;
3 organizing each computer storage disk into blocks; and
4 organizing the blocks into stripes, the stripes extending across a plurality of the
5 computer storage disks.

1 54. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 configuring the array as a RAID system of computer storage disks, the array hav-
3 ing n computer storage disks, and n is p+1, where p is a prime number;
4 organizing each computer storage disk into blocks;
5 organizing the blocks into stripes, the stripes extending across a plurality of the
6 computer storage disks.
7 forming n-2 rows of blocks, where each row contains one block from each disk.

1 55. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 using as the plurality of first storage devices units which have unequal lengths,
3 and each row has one row parity block and one diagonal parity block.

1 56. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 implementing control of a RAID organization of the array in logic circuitry.

1 57. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
2 using as the plurality of first storage devices units which are one of video tape,
3 magnetic tape, optical, DVD, bubble memory, magnetic disk, electronic random access
4 memory, and micro-electro mechanical storage devices.

- 1 58. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 using as the plurality of first storage devices media adapted to store information
 - 3 contained within the data and parity blocks.

- 1 59. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 using as the plurality of first storage devices media adapted to store information
 - 3 contained within the data and parity blocks.

- 1 60. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 storing diagonal parity blocks on the diagonal parity disk for each of the diagonal parity sets except one.

- 1 61. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:
 - 2 dividing each device into blocks; and
 - 3 organizing the blocks into stripes.

- 1 62. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:

2 configuring the plurality of first storage devices to refer to physical devices which
3 are not present, and physical devices which are not present are treated as containing zero
4 valued data.

1 63. (PREVIOUSLY PRESENTED) The method as in claim 42, further comprising:

2 providing the array with a predetermined number of storage devices equal to n,
3 where $n=p+1$ and p is a prime number, including n-2 data devices configured to hold
4 data, one row parity device and one diagonal parity device.

1 64. (PREVIOUSLY PRESENTED) A storage array organized to enable recovery from
2 two or fewer concurrent failures of storage devices, comprising:

3

4 means for storing data and row parity in stripes on a plurality of first storage devices, the
5 stripes configured as rows of blocks on the first devices;

6 means for defining diagonal parity sets that span the first devices;

7 means for computing a diagonal parity for the diagonal parity sets; and

8 means for storing the diagonal parity on a diagonal parity device, the diagonal
9 parity device separate from the first storage devices, the storage array including the plu-
10 rality of first storage devices and the diagonal parity device.

1 65. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for predetermining the a number of storage devices as n, where n is p+1,
3 and where p is a prime number.

1 66. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for storing row parity blocks on a single storage device of the plurality of
3 first storage devices.

1 67. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for configuring the plurality of first storage devices so that not all first
4 storage devices are present, and treating absent data storage devices as containing zero
5 valued data.

1 68. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for computing row parity blocks by an encoding of data blocks that is tol-
3 erant of single device failures.

1 69. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for configuring each row to have one row parity block and one diagonal
4 parity block.

1 70. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for placing parity blocks in different locations in different stripes from de-
4 vice to device.

1 71. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for implementing a row-diagonal (R-D) parity technique that computes the
4 diagonal parity along diagonal parity sets that span the first devices, and that stores the
5 diagonal parity for all of the diagonal parity sets except one on the diagonal parity device.

1 72. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for invoking storage access operations to and from the array in accordance
4 with a R-D parity technique.

1 73. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for storing row parity blocks computed for a stripe on a single first storage
3 device.

1 74. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for configuring the array as a RAID system of computer storage units.

1 75. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for configuring the array as a RAID system of computer storage disks;

3 means for organizing each computer storage disk into blocks; and

4 means for organizing the blocks into stripes, the stripes extending across a plural-
5 ity of the computer storage disks.

1 76. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for configuring the array as a RAID system of computer storage disks;

3 means for organizing each computer storage disk into blocks; and

4 means for organizing the blocks into stripes, the stripes extending across a plural-
5 ity of the computer storage disks.

6 means for forming n-2 rows of blocks, where each row contains one block from
7 each disk.

1 77. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for using as the plurality of first storage devices units which have unequal
3 lengths, and each row has one row parity block and one diagonal parity block.

1 78. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:
2 means for implementing control of a RAID organization of the array in logic cir-
3 cuitry.

1 79. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing
3 means for using as the plurality of first storage devices units which are one of
4 video tape, magnetic tape, optical, DVD, bubble memory, magnetic disk, electronic ran-
5 dom access memory, and micro-electro mechanical storage devices.

1 80. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for using as the plurality of first storage devices media adapted to store in-
4 formation contained within the data and parity blocks.

1 81. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further compris-
2 ing:

3 means for using as the plurality of first storage devices a plurality of data devices,
4 and the data devices are not all present, and absent data devices are treated as containing
5 zero valued data.

1 82. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:
2

3 means for storing diagonal parity blocks on the diagonal parity disk for each of
4 the diagonal parity sets except one.

1 83. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:

2 means for dividing each device into blocks; and

3 means for organizing the blocks into stripes.

1 84. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:
2

3 means for configuring the plurality of first storage devices to refer to physical devices which are not present, and physical devices which are not present are treated as
4 5 containing zero valued data.

1 85. (PREVIOUSLY PRESENTED) The storage array as in claim 64, further comprising:
2

3 means for providing the array with a predetermined number of storage devices
4 5 equal to n, where n=p+1 and p is a prime number, including n-2 data devices configured to hold data, one row parity device and one diagonal parity device.

1 86. (PREVIOUSLY PRESENTED) Electromagnetic Signals propagating on a computer
2 network, comprising:

3 said electromagnetic signals carrying instructions for execution on a processor for
4 the practice of a method for enabling recovery from two or fewer concurrent failures of
5 storage devices in a storage array, the method having the steps of,

6 storing data and row parity in stripes on a plurality of first storage devices, the
7 stripes configured as rows of blocks on the first devices;

8 defining diagonal parity sets that span the first devices;

9 computing a diagonal parity for the diagonal parity sets; and

10 storing the diagonal parity on a diagonal parity device, the diagonal parity device
11 separate from the first storage devices, the storage array including the plurality of first
12 storage devices and the diagonal parity device.

1 87. (PREVIOUSLY PRESENTED) Computer readable media, comprising:

2 said computer readable media having instructions written thereon for execution on
3 a processor for the practice of a method for enabling recovery from two or fewer concurrent
4 failures of storage devices in a storage array, the method having the steps of,

5 storing data and row parity in stripes on a plurality of first storage devices, the
6 stripes configured as rows of blocks on the first devices;

7 defining diagonal parity sets that span the first devices;

8 computing a diagonal parity for the diagonal parity sets; and

9 storing the diagonal parity on a diagonal parity device, the diagonal parity device
10 separate from the first storage devices, the storage array including the plurality of first
11 storage devices and the diagonal parity device.